arXiv:1512.05840v1 [cs.CY] 18 Dec 2015

Deep Poisson Factorization Machines: factor analysis for mapping behaviors in journalist ecosystem

Pau Perng-Hwa Kung
Media Lab
Massachusetts Institute of Technology
Cambridge, MA 01239
pernghwa@media.mit.edu

Abstract

Newsroom in online ecosystem is difficult to untangle. With prevalence of social media, interactions between journalists and individuals become visible, but lack of understanding to inner processing of information feedback loop in public sphere leave most journalists baffled. Can we provide an organized view to characterize journalist behaviors on individual level to know better of the ecosystem? To this end, I propose Poisson Factorization Machine (PFM), a Bayesian analogue to matrix factorization that assumes Poisson distribution for generative process. The model generalizes recent studies on Poisson Matrix Factorization to account temporal interaction which involves tensor-like structure, and label information. Two inference procedures are designed, one based on batch variational EM and another stochastic variational inference scheme that efficiently scales with data size. An important novelty in this note is that I show how to stack layers of PFM to introduce a deep architecture. This work discusses some potential results applying the model and explains how such latent factors may be useful for analyzing latent behaviors for data exploration.

1 Intuition and Introduction

Digital newsroom is rapidly changing. With the introduction of social media, the new public sphere offers greater mutual visibility and diversity. However, a full understanding of how the new system operates is still out of complete grasp. In particular, the problem is made further challenging with data sparsity, where individual records about responses to any particular news event is limited.

We believe mapping the development of analytical framework that integrates different domains of data can significantly help us characterize behavioral patterns of journalists and understand how they operate in social media according to different news topics and events. Figure 1 outlines a possible data mapping. Ingesting data from published feeds by news organizations and map corresponding journalists to social media and observe their activities will undoubtedly enable a clearer mapping of individual behaviors related to news event progression. Based on this data mapping, we can incorporate general response by normal users on social media to construct a data feedback loop.

We are interested in aggregated behaviors as well as individual behaviors that interact with the ecosystem. A way to achieve this aim is to design a mathematical model that can map the latent process that characterize individual behavior as well as feature analysis on aggregate level. In this work, we propose a model building on top of probabilistic matrix factorization, called Poisson Factorization Machine, to describe the latent factors that behavioral features are embedded upon. In particular, we first create a data design matrix by inferring links between data instances from different domains. Figure 2 illustrates such a representation, where each row is a news article published by news organization. The columns denote the features associated with the story, which can include

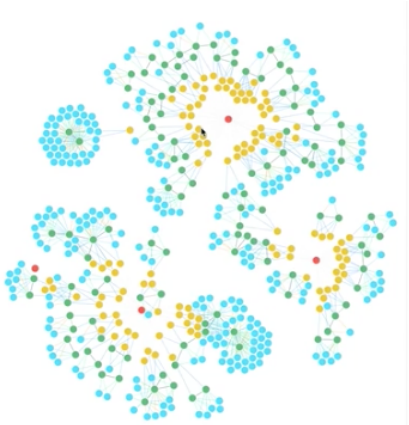

Figure 1: A sampled data view of a network of relational data spanning news data and social media data. Blue nodes denote tweets, green nodes denote journalists, yellow nodes denote news articles, and red nodes denote news organizations. Edges are data relations and interactions.

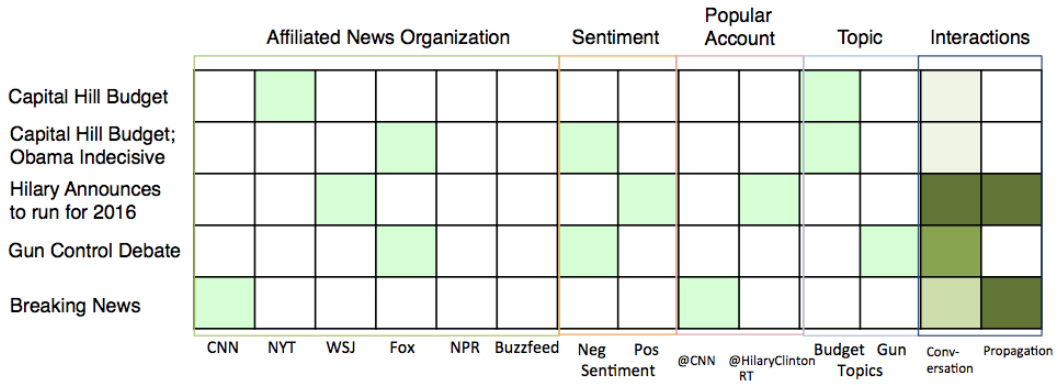

Figure 2: An example encoded under factorization machine format. Each row encodes an item with corresponding features. A color box denotes a particular feature category. Color shades denotes frequency; darker means higher counts. Different categories can interact and particular category of feature encodes dependency. We can dictate certain categories as response variables for supervision.

topics written in the news, name entities mentioned, as well as social media engagement associated with journalists' Twitter accounts and general users interactions. Then, we perform our Poisson Factorization Machine to learn a lower dimensional embedding of the latent factors that map to these features.

Using this framework, we can answer several key questions: **query-** given any data instance or any subset of features as queries, can we understand potentially related instances and features by the learned latent feature mapping? **factor analysis-** can we describe the learned latent factors in a meaningful way to better understand relations between features? **prediction-** given new data instances, can we use its latent dimensional embedding to predict the instance's response at some unobserved features (e.g. features that have not yet occurred so we cannot record).

Our proposed factorization model offers several advantages. First, the model bases on a hierarchical generative structure with Poisson distribution, which provide better statistical strength sharing among factors and model human behavior data better, being discrete output. Second, our model generalizes earlier works on Poisson Matrix Factorization by extending the model to include response variables to enable a form of supervised guidance. We offer batch as well as online variational inference to accommodate large-scale datasets. In addition, we show that it is possible to stack units of PFM to make a deep architecture. Third, we further extend our model to incorporate multi-way

interactions among features. This provides natural feature grouping for more interpretable feature exploration.

2 Methodology

In this section, we describe the details for the proposed Poisson Factorization Machine, where we discuss the formal definition of the model, and inference with prediction procedures.

Consider a set of observed news articles, where for each article, we have mapped the authors to their respective Twitter handles and extracted a set of behavior patterns around the time of posting the article. We denote this data matrix by $\mathcal{X} \in \mathbb{R}^{NxD}$, where each entry x_i is a D dimensional feature associated with that article. Also, we have response variables $\mathcal{Y} \in \mathbb{R}^{NxC}$ with total of C response types. In this work, we set C = 1, but the reasoning can easily apply to the higher dimension cases.

Imagine there are $k \in \{1,\cdots,K\}$ latent factors from the matrix factorization model. We define the Poisson Factorization Machine with the following generative process:

$$\begin{aligned}
\theta_{i,k} &\sim Gamma(a, ac) \\
\beta_{k,d} &\sim Gamma(b, b) \\
x_{i,d} &\sim Poisson(\textstyle\sum_{k=1}^{K} \theta_{i,k}\beta_{k,d}) \\
y_{i,c} &\sim Gaussian(\theta_i^T \eta, \sigma)
\end{aligned}$$

where $\beta_k \in \mathbb{R}_+^D$ is the kth latent factor mapping for feature vector, $\theta_i \in \mathbb{R}_+^K$ denotes the factor weights. Note the hyperparameters of the model a, b, c, and $\eta \in \mathbb{R}^{KxC}, \sigma \in \mathbb{R}^{CxC}$. Here, c is a scalar used to tune the weights. To generate the response variable, for each instance, we use the corresponding factor weights as input for normal linear model for regression. Figure shows the graphical model. In this work, we leave a, b predefined, and update other hyperparameters during model fitting.

Notice the introduction of auxiliary variable z. This is employed to facilitate posterior variational inference by helping us to develop closed coordinate ascent updates for the parameters. More specifically, by the Superposition Theorem [4], the Poisson distribution dictated by x can be composed of K independent Poisson variables $x_{id} = \Sigma_{k=1}^{K} z_{idk}$ where $z_{idk} \sim Poisson(\theta_{ik}\beta_{kd})$. Effectively, when we are looking at the data likelihood, it can be viewed as marginalizing z (if we denote all hidden parameters by π), but keep in mind that having z not marginalized maintains flexibility for inference later on:

$$P(y_{i,c}, x_{i,d}|\pi) = \sum_{z_{idk}} P(y_{i,c}, x_{i,d}, z_{i,d,k}|\pi_{-z})$$

The intuition for selecting Poisson model comes from the statistical property that Poisson distribution models count data, which is more suitable for sparse behavioral patterns than traditional Gaussian-based matrix factorization. Also, Poisson factorization avoids the problem of data sparsity since it penalizes 0 less strongly than Gaussian distribution, which makes it desirable for weakly labeled data[1]. The selection of Gamma prior as originally suggested by [3][5]. Notice that response variable need not be Gaussian, and can vary depending on the nature of inputting dataset. In fact, since our proposed learning and inference procedure is very similar as in [1], our model can generalize to Generalized Linear Models, which include other distributions such as Poisson, Multinomial, etc.

2.1 Posterior Inference

For posterior inference, we adopt variational EM algorithm, which was employed in Supervised Topic Model [1]. After showing the derivation, we will also present a stochastic variational inference analogue for the inference algorithm.

First, we consider a fully factorized variational distribution for the parameters:

$$q(\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{z}) = \prod_{k=1}^{K} \Big(\prod_{i=1}^{N} q(\theta_{ik}) \prod_{d=1}^{D} q(z_{idk}) \Big) \Big(\prod_{d=1}^{D} q(\beta_{kd}) \Big)$$

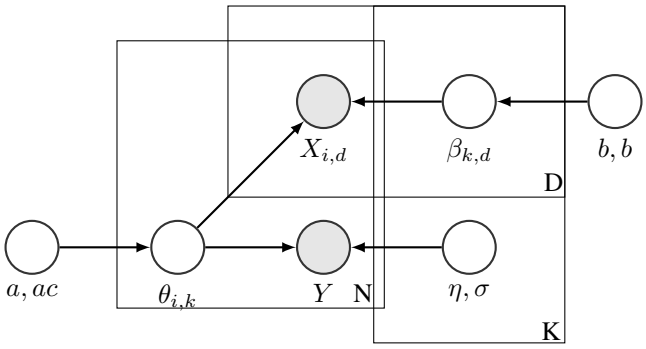

Figure 3: The graphical model representation of Poisson Factorization Model (PFM). Shaded nodes are observed variables, and note that we did not include auxiliary variable z.

For the above variational distribution, we choose the following distributions to reflect the generative process in the Poisson Factorization Model:

$$\begin{aligned} q(\theta_{i,k}) &= Gamma(\gamma_{ik}, \chi_{ik}) \\ q(\beta_{k,d}) &= Gamma(\nu_{kd}, \lambda_{kd}) \\ q(z_{i,k,d}) &= Multinomial(\overrightarrow{\phi_{id}}) \end{aligned}$$

The log likelihood for the data features is calculated and bounded by:

$$\begin{aligned} ln\ p(x_{1:N}|\pi) &= ln \int_\theta \int_\beta p(\theta, \beta, x, y|\pi_{-\theta,-\beta}) \frac{q(\theta,\beta)}{q(\theta,\beta)} \\ &\geq \mathbb{E}_{q}[ln\ p(\theta, \beta, x|\pi)] - \mathbb{E}_{q}[ln\ q(\theta, \beta)] \\ &= \mathbb{E}_{q}[ln\ p(\theta, \beta, x|\pi)] + H(q) \end{aligned}$$

where $H(q)$ is the entropy of the variational distribution q and the resulting lower bound is our Evidence Lower Bound (ELBO).

2.1.1 Batch inference

The ELBO lower bounds the log-likelihood for all article items in the data. For the E-step in variational EM procedure, we run variational inference on the variational parameters. In the M-step, we optimize according to the ELBO with respect to model parameters c (with a, b held fixed), as well as regression parameters η and σ.

Variational E-step To calculate the updates, we try to first further break down $E_q[ln\ p(\theta, \beta, x|\pi)]$ and draw analogy to earlier works on Poisson Matrix Factorization [2][3][5][7] to illustrate how we can directly leverage on previous results to derive coordinate ascent update rules for the proposed model. Notice that:

$$\begin{aligned} \mathbb{E}_{q}[ln\ p(\theta, \beta, x|\pi)] = \mathbb{E}_{q}[ln\ p(\beta|b, b)] \\ + \mathbb{E}_{q}[ln\ p(\theta|a, ac)] + \sum_{n=1}^{N} \mathbb{E}_{q}[ln\ p(z|x, \theta, \beta)] \\ + \sum_{n=1}^{N} \mathbb{E}_{q}[ln\ p(x|\theta, \beta)] + \mathbb{E}_{q}[ln\ p(y|\theta, \eta, \sigma)] \end{aligned}$$

In fact, for variational parameters, when taking partial derivatives, the expression is identical to batch inference algorithm in ordinary Poisson Matrix Factorization, so we arrive at the following updates for $\theta_{i,k},\ \beta_{k,d}$:

$$\overrightarrow{\phi_{idk}} := \frac{exp(\mathbb{E}_{q}[ln\ \theta_{ik}\beta_{kd}])}{\sum_{k=1}^{K} exp(\mathbb{E}_{q}[ln\ \theta_{ik}\beta_{kd}])}$$

$$\gamma_{i,k} = a + \sum_{d=1}^{D} x_{id}\phi_{idk}$$
$$\chi_{i,k} = ac + \sum_{d=1}^{D} \mathbb{E}_{\text{ıı}}[\beta_{kd}]$$
$$\nu_{k,d} = b + \sum_{i=1}^{N} x_{id}\phi_{idk}$$
$$\lambda_{k,d} = b + \sum_{i=1}^{N} \mathbb{E}_{\text{ıı}}[\theta_{ik}]$$

Expectations for $\mathbb{E}_{\text{ıı}}[\theta_{ik}]$ and $\mathbb{E}_{\text{ıı}}[ln\ \theta_{ik}]$ are γ_{ik}/χ_{ik} and $\psi(\gamma_{ik}) - ln\ \chi_{ik}$ respectively (where $\psi(\cdot)$ is the digamma function).

Variational M-step In the M-step, we maximize the article-level ELBO with respect to c, η, and σ.

We update c according to

$$c = \mathbb{E}_{\text{ıı}}[\bar{\theta}] \quad \Rightarrow \quad c^{-1} = \frac{1}{NK}\sum_{i,k}\mathbb{E}_{\text{ıı}}[\theta_{ik}].$$

For regression parameters, we take partial derivatives with respect to η and σ, yielding for η:

$$\frac{\partial \mathcal{L}}{\partial \eta} = \frac{\partial}{\partial \eta}\Big(\frac{1}{\sigma}\Big)\sum_{i=1}^{N}\mathbb{E}[\eta^T\theta_i]y_i - \mathbb{E}[\frac{\eta^T\theta_i\theta_i^T\eta_i}{2}]$$
$$= \frac{1}{\sigma}\sum_{i=1}^{N}\mathbb{E}[\theta_i]y_i - \eta^T\mathbb{E}[\theta_i\theta_i^T]$$

where the latter term $\mathbb{E}[\theta_i\theta_i^T]$ is evaluated as

$$\frac{1}{K^2}\Big(\sum_j\sum_{l\neq j}\frac{\gamma_{ij}}{\chi_{ij}}\frac{\gamma_{il}}{\chi_{il}} + \sum_j(\frac{1}{ac^2} + \frac{\gamma_{ij}}{\chi_{ij}})\Big).$$

Setting the partial derivative to zero and we arrive at the updating rule for η, as we collapse all instances i into vector:

$$\eta = \Big(\mathbb{E}[\theta^T\theta]\Big)^{-1}\mathbb{E}[\theta]^T y$$

For derivative of σ, we have:

$$\frac{\partial \mathcal{L}}{\partial \sigma} = \frac{-1}{2\sigma} + \Big(\frac{1}{\sigma}\Big)\sum_{i=1}^{N}\mathbb{E}[\eta^T\theta_i]y_i - \mathbb{E}[\frac{\eta^T\theta_i\theta_i^T\eta_i}{2}] \Rightarrow$$
$$\sigma = \frac{1}{N}\Big\{y^Ty - y^T\mathbb{E}[\theta]\Big(\mathbb{E}[\theta^T\theta]\Big)^{-1}\mathbb{E}[\theta]^Ty\Big\}$$

To complete the procedure, we iterate between updating and optimizing parameters by following E-step and M-step until convergence. Algorithm 1 outlines the entire process.

2.1.2 Stochastic Inference

Batch variational EM algorithm might take multiple passes at the dataset before reaching convergence. In real world applications, this may not be desirable since social media data is more dynamic and large in scale. Recently, Hoffman et al. developed stochastic variational inference where each iteration we work with a subset of data instead, averaging over the noisy variational objective $\mathcal{L}$. We optimize by iterating between global variables and local variables. The global objective is based on β, and at each iteration t, with data subset B_t,we have:

$$\mathcal{L}_t = \frac{N}{|B_t|}\sum_{i\in B_t}\mathbb{E}_q[ln\ p(x_i,\theta_i|\beta)] + \mathbb{E}_q[ln\ p(\beta)] + H(q).$$

The main difference in the inference procedure for stochastic variational inference is in the updating part, where parameters c, η, and σ are updated with the average empirically according to the mini-batch samples (where $\widehat{\theta} = \{\theta_i;\ \forall i \in B_t\}$, $\widehat{y} = \{y_i;\ \forall i \in B_t\}$):

$$(c^{(t)})^{-1} = \frac{1}{|B_t|K} \sum_{i \in B_t, k} \mathbb{E}_{\shortparallel}[\theta_{ik}]$$

$$\eta^{(t)} = \Big(\mathbb{E}[\widehat{\theta}^T \widehat{\theta}]\Big)^{-1} \mathbb{E}[\widehat{\theta}]^T \widehat{y}$$

$$\sigma^{(t)} = \frac{1}{|B_t|}\Big\{\widehat{y}^T \widehat{y} - \widehat{y}^T \mathbb{E}[\widehat{\theta}]\Big(\mathbb{E}[\widehat{\theta}^T \widehat{\theta}]\Big)^{-1} \mathbb{E}[\widehat{\theta}]^T \widehat{y}\Big\}$$

As for β, since it is also global variable, we update by taking a natural gradient step based on the Fisher information matrix of variational distribution $q(\beta_{kd})$. We have the variational parameters for β updated as:

$$\nu_{k,d}^{(t)} = (1 - \rho_t)\nu_{kd}^{(t-1)} + \rho_t\Big(b + \frac{N}{|B_t|} \sum_{i \in B_t} x_{id}\phi_{idk}\Big)$$

$$\lambda_{k,d}^{(t)} = (1 - \rho_t)\lambda_{kd}^{(t-1)} + \rho_t\Big(b + \frac{N}{|B_t|} \sum_{i \in B_t} \mathbb{E}_{\shortparallel}[\theta_{ik}]\Big)$$

where ρ_t is a positive interpolation factor and ensuring convergence. We adopt $\rho_t = (t_0 + t)^{-\kappa}$ for $t_0 > 0$ and $\kappa \in (0.5, 1]$ which is shown to correspond to natural gradient step in information geometry with stochastic optimization[3][5]. Algorithm 2 summaries the stochastic variational inference procedure.

2.2 Query and Prediction

For any given query x_i or $s \subseteq D$, we can infer the corresponding latent factor distribution by computing $\mathbb{E}[\theta_i]$ as well as $\mathbb{E}[\beta_{-s}|s]$ given that $\mathbb{E}[\beta]$ is already learned by the model. Prediction follows by combining both latent factors to generate the response.

2.3 Modeling Cross-Category Interactions

Notice that we can enforce each $\theta_{i,k},\ \beta k, d$ to be an f dimensional vector. In this way, we create an f-way interaction among latent factors. We can further impose a sparse deterministic prior **B** with binary entries where each data entry encodes whether the two latent factors should interact. We can also approach by using a diffusion tree prior, which consider a hierarchy of prior interactions. Under this construct, we can flatmap the tensor data structure to two dimensional structure and apply the factor interactions to encode a higher level of feature dependency.

3 Deep Poisson Factorization Machines

Given the development of Poisson Factorization Machines, here we show that it is possible to stack layers of latent Gamma units on top of PFM to introduce a deep architecture. In particular, for latent factors β and θ, a layering (for layer l) parent variables $z_{n,l}$ with corresponding weights $w_{l,k}$ are fed with link function g_{l-1} specified by $g_{l-1}(z_l^T w_{l,k})$. As described in the recent work on Deep Exponential Families [6] with exponential family in the form

$$p(x) = h(x)\ exp(\eta^T T(x) - a(\eta))$$

, we can condition the sufficient statistics, the mean to be particular, controlled by the link function g_l via the gradient of the log-normalizer.

$$E[T(z_{l,k})] = \bigtriangledown_\eta a(g_l(z_{l+1}^T w_{l,k})).$$

Figure 4 shows the schematics of stacking layers of latent variables on PFM latent states.

Gamma variables The Gamma distribution is an exponential family distribution with the probability density controlled by natural parameters $\alpha,\ \beta$ where

$$p(z) = z^{-1} exp(\alpha log(z) - \beta z - log\Gamma(\alpha) - \alpha log(\beta))$$

Algorithm 1 Variational Inference for PFM

repeat
  Determine instance set $\mathcal{S}$ for iteration:
  if $batch\ inference$ **then**
    $\mathcal{S} = X \in \mathbb{R}^{NxD}$
  else
    $\mathcal{S} = B_t \in \mathbb{R}^{|B_t|xD}$
  end if
  E-Step:
  for each instance i **do**
    update $\gamma_i,\ \chi_i$
  end for
  for each component k **do**
    if $batch\ inference$ **then**
      update $\nu_k,\ \lambda_k$
    else
      update $\nu_k^{(t)},\ \lambda_k^{(t)}$
    end if
  end for
  update all corresponding ϕ_{idk}
  M-Step:
  if $batch\ inference$ **then**
    update $c,\ \eta,\ \sigma$
  else
    update $c^{(t)},\ \eta^{(t)},\ \sigma^{(t)}$
  end if
until convergence

The link functions for α and β are respectively

$$g_\alpha = \alpha_l,\ g_\beta = \frac{\alpha_l}{z_{l+1}^T w_{l,k}}$$

where α_l denotes the shape parameter at layer l, and $\mathbf{w}$ is simply drawn from a Gamma distribution.

Gaussian variables For Gaussian distribution, we are interested in the switching variables θ, which is calculated as part of stacking Gamma variables. **Inference** For inference, since all variables in

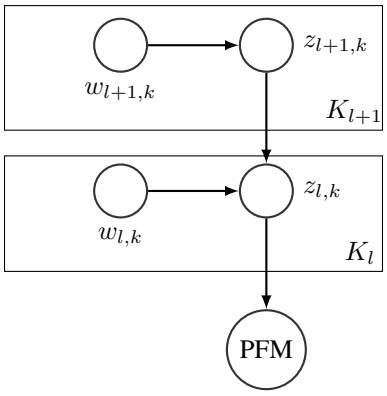

Figure 4: Schematics for stacking gamma layers on PFM.

each layer follows precisely for the inference procedures derived earlier, we focus on the stacking variables. The ELBO for the generative process now includes layer-wise w and z. For each stacking

z and corresponding w, we have:

$$w_l \sim q(w_l)$$
$$z_l \sim q(z_l)$$
$$g_l \sim \bigtriangledown log\ q(z_l)$$

For parameter updates, taking the gradient for the variational bound $\bigtriangledown_{z,w}\mathcal{L}$ can be used to update each stacking layer's parameters, respectively.

References

[1] D. Blei and J. McAuliffe. Supervised topic models. In NIPS, 2007.

[2] Z. Gan, C. Chen, R. Henao, D. Carlson, and L. Carin. Scalable deep poisson factor analysis for topic modeling. In ICML, 2015.

[3] P. Gopalan, J. Hofman, and D. Blei. Scalable Recommendation with Hierarchical Poisson Factorization, UAI, 2015.

[4] Kingman, John Frank (1992). Poisson processes. Claredon Press.

[5] D. Liang, J. Paisley, and D. Ellis. Codebook-based scalable music tagging with Poisson matrix factorization. In Proceedings of the Fifteenth International Society for Music Information Retrieval Conference, 2015.

[6] Ranganath, R., Tang, L., Charlin, L., and Blei, D. M. Deep exponential families. AISTATS, 2015.

[7] A. Schein, J. Paisley, D. M. Blei, and H. Wallach. Bayesian poisson tensor factorization for inferring multilateral relations from sparse dyadic event counts. SIGKDD, 2015.